

WHAT IS CLAIMED IS:

- 1 1. A method for sending data from a source to a destination, comprising:
2 a host providing to a sending agent of the source, virtual memory addresses of data to
3 be sent to a destination wherein the data is stored in a plurality of physical locations of the
4 source, each location having a physical address and a virtual memory address which is mapped
5 to the physical address;
6 the sending agent providing to the host at least some of the virtual memory addresses of
7 the data to be sent to the destination;
8 the host identifying to the sending agent the data addressed by the virtual memory
9 addresses provided by the sending agent; and
10 the sending agent sending the identified data to the destination.

- 1 2. The method of claim 1 wherein the host identifying data comprises the host
2 providing to the sending agent the data addressed by the virtual addresses provided by the
3 sending agent, said method further comprising the sending agent storing the data received from
4 the host in a buffer of the sending agent.

- 1 3. The method of claim 1 wherein the host identifying data comprises the host
2 providing to the sending agent the physical addresses of the locations containing the data
3 addressed by the virtual memory addresses provided by the sending agent.

- 1 4. The method of claim 3 wherein the physical locations include locations of a first
2 memory and locations of a second memory and the data identified by the host is stored in the
3 first memory and the physical memory addresses provided by the host are physical memory
4 locations of the first memory containing the data addressed by the virtual addresses provided by
5 the sending agent, the method further comprising pinning the physical memory locations of the

6 first memory provided by the host to the sending agent to prevent the data addressed by the
7 virtual addresses provided by the sending agent from being swapped to the second memory.

1 5. The method of claim 4 further comprising:
2 the sending agent retrieving from the pinned physical memory locations of the first
3 memory, the data addressed by the virtual addresses provided by the sending agent; and
4 unpinning the pinned physical memory locations of the first memory after the sending
5 agent sends to the destination the data addressed by the virtual addresses provided by the
6 sending agent.

1 6. The method of claim 1 further comprising receiving from the destination an
2 acknowledgment for data successfully sent by the sending agent and received by the
3 destination; wherein the virtual memory addresses provided by the sending agent to the host are
4 the virtual addresses of data sent by the sending agent to the destination but not acknowledged
5 as successfully received by the destination.

1 7. The method of claim 1 further comprising:
2 receiving from the destination an acknowledgment for data successfully sent by the
3 sending agent and received by the destination;
4 the sending agent providing to the host the virtual addresses of data sent by the sending
5 agent to the destination but not acknowledged as successfully received by the destination;
6 the host identifying to the sending agent the unacknowledged data addressed by the
7 virtual memory addresses provided by the sending agent; and
8 the sending agent resending the identified unacknowledged data to the destination.

1 8. The method of claim 1 wherein the host providing virtual addresses to the
2 sending agent includes the host providing to the sending agent at least one data structure which
3 includes in an address field containing the virtual address of one of a plurality of memory
4 locations storing a block of data to be sent to the destination, a size field containing a value
5 representing the size of the block of data; and a sequence number field containing a value
6 representing a packet sequence number associated with data within the block of data.

1 9. The method of claim 1 wherein the host providing virtual addresses to the
2 sending agent includes the host providing to the sending agent a plurality of data structures,
3 wherein each data structure includes in an address field, the virtual address of one of a plurality
4 of memory locations storing a block of data to be sent to the destination, a size field containing
5 a value representing the size of the block of data, a sequence number field containing a value
6 representing the packet sequence number associated with data within the block of data; and a
7 link field containing the virtual address of another data structure of the plurality of data
8 structures.

1 10. The method of claim 1 wherein the physical locations include locations of a first
2 memory and a second memory and the data to be sent to the destination is stored in the first
3 memory, the method further comprising,
4 pinning the locations of the first memory storing the data to be sent to prevent the data
5 to be sent from being swapped to the second memory;
6 the host providing to the sending agent in addition to the virtual memory addresses of
7 the data to be sent, the physical addresses of the locations of the first memory storing the data
8 to be sent;

9 the sending agent retrieving from the pinned locations of the first memory, the data to
10 be sent; and
11 unpinning the pinned locations of the first memory storing the data to be sent after the
12 sending agent retrieves the data from the pinned locations of the first memory storing the data to
13 be sent.

1 11. A system adapted to communicate with a destination, comprising:
2 memory;
3 a processor coupled to the system memory;
4 an operating system executable by the processor in memory;
5 a network adaptor;
6 data storage;
7 a data storage controller adapted to manage Input/Output (I/O) access to the data
8 storage; and
9 a device driver executable by the processor in the memory,
10 wherein the memory and the data storage each comprise physical locations adapted to
11 store data, each location having a physical address and a virtual address which is mapped to
12 the physical address; and
13 wherein at least one of the operating system and device driver is adapted to provide a
14 host and at least one of the device driver and the network adaptor is adapted to provide a
15 sending agent wherein:
16 (I) the host provides to the sending agent, virtual memory addresses of data to
17 be sent to a destination,
18 (ii) the sending agent provides to the host at least some of the virtual memory
19 addresses of the data to be sent to the destination;

20 (iii) the host identifies to the sending agent the data addressed by the virtual
21 memory addresses provided by the sending agent; and

22

23

24

25 (iv) the sending agent sending the identified data to the destination.

26 establish an active connection adapted to send packets of data between the system and
a destination.

1 12. The system of claim 11 wherein the system further comprises a buffer and
2 wherein the host identifying data comprises the host providing to the sending agent the data
3 addressed by the virtual addresses provided by the sending agent, and wherein the sending
4 agent is further adapted to store the data received from the host in the buffer.

1 13. The system of claim 11 wherein the host identifying data comprises the host
2 providing to the sending agent the physical addresses of the locations containing the data
3 addressed by the virtual memory addresses provided by the sending agent.

1 14. The system of claim 13 wherein the data identified by the host is stored in the
2 memory and the physical addresses provided by the host are physical locations of the memory
3 containing the data addressed by the virtual addresses provided by the sending agent, the
4 system further comprising pinning the physical memory locations of the memory provided by the
5 host to the sending agent to prevent the data addressed by the virtual addresses provided by
6 the sending agent from being swapped to the data storage.

1 15. The system of claim 14 wherein the sending agent is further adapted to retrieve
2 from the pinned physical memory locations of the memory, the data addressed by the virtual
3 addresses provided by the sending agent; and
4 at least one of the sending agent and the host is further adapted to unpin the pinned
5 physical memory locations of the memory after the sending agent sends to the destination the
6 data addressed by the virtual addresses provided by the sending agent.

1 16. The system of claim 11 wherein the sending agent is further adapted to receive
2 from the destination an acknowledgment for data successfully sent by the sending agent and
3 received by the destination; and wherein the virtual memory addresses provided by the sending
4 agent to the host are the virtual addresses of data sent by the sending agent to the destination
5 but not acknowledged as successfully received by the destination.

1 17. The system of claim 11 wherein the sending agent is further adapted to:
2 receive from the destination an acknowledgment for data successfully sent by the
3 sending agent and received by the destination; and
4 provide to the host the virtual addresses of data sent by the sending agent to the
5 destination but not acknowledged as successfully received by the destination;
6 wherein the host is further adapted to identify to the sending agent the unacknowledged
7 data addressed by the virtual memory addresses provided by the sending agent; and
8 wherein the sending agent is further adapted to resend the identified unacknowledged
9 data to the destination.

1 18. The system of claim 11 wherein the host providing virtual addresses to the
2 sending agent includes the host providing to the sending agent at least one data structure which
3 includes an address field containing the virtual address of one of a plurality of locations storing a

4 block of data to be sent to the destination, a size field containing a value representing the size of
5 the block of data; and a sequence number field containing a value representing a packet
6 sequence number associated with data within the block of data.

1 19. The system of claim 11 wherein the host providing virtual addresses to the
2 sending agent includes the host providing to the sending agent a plurality of data structures,
3 wherein each data structure includes an address field containing the virtual address of one of a
4 plurality of memory locations storing a block of data to be sent to the destination, a size field
5 containing a value representing the size of the block of data, a sequence number field containing
6 a value representing the packet sequence number associated with data within the block of
7 data; and a link field containing the virtual address of another data structure of the plurality of
8 data structures.

1 20. The system of claim 11 wherein the data to be sent to the destination is stored
2 in the memory, and wherein at least one of the host and the sending agent is adapted to:
3 pin the locations of the memory storing the data to be sent to prevent the data to be
4 sent from being swapped to the data storage;
5 wherein the host is further adapted to provide to the sending agent in addition to the
6 virtual memory addresses of the data to be sent, the physical addresses of the locations of the
7 memory storing the data to be sent;
8 wherein the sending agent is further adapted to retrieve from the pinned locations of the
9 memory, the data to be sent; and to unpin the pinned locations of the memory storing the data
10 to be sent after the sending agent retrieves the data from the pinned locations of the memory
11 storing the data to be sent.

1 21. An article of manufacture for sending data from a source to a destination, the
2 operations comprising:
3 a host providing to a sending agent of the source, virtual memory addresses of data to
4 be sent to a destination wherein the data is stored in a plurality of physical locations of the
5 source, each location having a physical address and a virtual memory address which is mapped
6 to the physical address;
7 the sending agent providing to the host at least some of the virtual memory addresses of
8 the data to be sent to the destination;
9 the host identifying to the sending agent the data addressed by the virtual memory
10 addresses provided by the sending agent; and
11 the sending agent sending the identified data to the destination.

1 22. The article of manufacture of claim 21 wherein the host identifying data
2 comprises the host providing to the sending agent the data addressed by the virtual addresses
3 provided by the sending agent, said operations further comprising the sending agent storing the
4 data received from the host in a buffer of the sending agent.

1 23. The article of manufacture of claim 21 wherein the host identifying data
2 comprises the host providing to the sending agent the physical addresses of the locations
3 containing the data addressed by the virtual memory addresses provided by the sending agent.

1 24. The article of manufacture of claim 23 wherein the physical locations include
2 locations of a first memory and locations of a second memory and the data identified by the
3 host is stored in the first memory and the physical memory addresses provided by the host are
4 physical memory locations of the first memory containing the data addressed by the virtual
5 addresses provided by the sending agent, the operations further comprising pinning the physical
6 memory locations of the first memory provided by the host to the sending agent to prevent the

7 data addressed by the virtual addresses provided by the sending agent from being swapped to
8 the second memory.

1 25. The article of manufacture of claim 24 wherein the operations further comprise:
2 the sending agent retrieving from the pinned physical memory locations of the first
3 memory, the data addressed by the virtual addresses provided by the sending agent; and
4 unpinning the pinned physical memory locations of the first memory after the sending
5 agent sends to the destination the data addressed by the virtual addresses provided by the
6 sending agent.

1 26. The article of manufacture of claim 21 wherein the operations further comprise
2 receiving from the destination an acknowledgment for data successfully sent by the sending
3 agent and received by the destination; wherein the virtual memory addresses provided by the
4 sending agent to the host are the virtual addresses of data sent by the sending agent to the
5 destination but not acknowledged as successfully received by the destination.

1 27. The article of manufacture of claim 21 wherein the operations further comprise:
2 receiving from the destination an acknowledgment for data successfully sent by the
3 sending agent and received by the destination;
4 the sending agent providing to the host the virtual addresses of data sent by the sending
5 agent to the destination but not acknowledged as successfully received by the destination;
6 the host identifying to the sending agent the unacknowledged data addressed by the
7 virtual memory addresses provided by the sending agent; and
8 the sending agent resending the identified unacknowledged data to the destination.

1 28. The article of manufacture of claim 21 wherein the host providing virtual
2 addresses to the sending agent includes the host providing to the sending agent at least one data
3 structure which includes in an address field containing the virtual address of one of a plurality of
4 memory locations storing a block of data to be sent to the destination, a size field containing a
5 value representing the size of the block of data; and a sequence number field containing a value
6 representing a packet sequence number associated with data within the block of data.

1 29. The article of manufacture of claim 21 wherein the host providing virtual
2 addresses to the sending agent includes the host providing to the sending agent a plurality of
3 data structures, wherein each data structure includes in an address field, the virtual address of
4 one of a plurality of memory locations storing a block of data to be sent to the destination, a
5 size field containing a value representing the size of the block of data, a sequence number field
6 containing a value representing the packet sequence number associated with data within the
7 block of data; and a link field containing the virtual address of another data structure of the
8 plurality of data structures.

1 30. The article of manufacture of claim 21 wherein the physical locations include
2 locations of a first memory and a second memory and the data to be sent to the destination is
3 stored in the first memory, the operations further comprising,
4 pinning the locations of the first memory storing the data to be sent to prevent the data
5 to be sent from being swapped to the second memory;
6 the host providing to the sending agent in addition to the virtual memory addresses of
7 the data to be sent, the physical addresses of the locations of the first memory storing the data
8 to be sent;

9 the sending agent retrieving from the pinned locations of the first memory, the data to
10 be sent; and
11 unpinning the pinned locations of the first memory storing the data to be sent after the
12 sending agent retrieves the data from the pinned locations of the first memory storing the data
13 to be sent.